Research engineer position (6 months, potentially extensible), to start as soon as possible

Title: Learning curiosity-driven question generation and evaluation with deep language models (towards educational technology applications)

Supervision: Pierre-Yves Oudeyer (Inria and Microsoft Research Montreal), Eric Yuan and Tong Wang (Microsoft Research Montreal), Hélène Sauzéon (Inria)

Host team: Flowers team, Inria Bordeaux, https://flowers.inria.fr/ and involves a collaboration with MSR Montreal.

Duration: 6 months, starting from April, 1st 2022 (potentially extensible)

Notes: We will be prioritizing candidates already in France due to administrative reasons and need to start fast this project.

How to apply: contact t-poudeyer@microsoft.com, and helene.sauzeon@inria.fr (send one email to all) with a CV and letter of motivation

Keywords: question generation; question evaluation; curiosity-driven learning; NLP; deep language models; educational technology; AI for humanity

Context:

This project is a machine learning project, in the domain of natural language processing and using deep learning techniques, studying algorithms enabling AI systems to ask curiosity-driven questions given a piece of text. It is happening in an educational technology research context aiming to provide children with tools that foster development of curiosity-driven learning skills: the outcomes of this project will be reused in further educational technology experiments.

Educational technology context: Curiosity is an important factor that enables personalized and independent learning in children. Research suggests that it is a skill that might be practiced by promoting information-searching behaviors. In an ongoing project, the Flowers lab has been studying how a conversational agent can foster question-asking skills and develop curiosity in primary school children (Alaimi et al., 2020; Abdelghani et al. 2021). This is of high educational and societal importance, as most occidental educational systems tend to work against the development of curiosity in classrooms. However, research on curiosity-driven

learning shows 1) the importance and efficiency of this form of learning 2) curiosity happens when one becomes aware of what one does not know (transforming unknown unknown into known unknown); 3) this meta-cognitive skills is strongly linked with the ability to ask divergent questions about a topic (a divergent question is a question relevant to the topic but for which the answer is not already known/in the test/video giving information about the topic).

In a recent paper (Alaimi et al., 2020) and ongoing work (Abdelghani et al., 2021), we showed that certain forms of linguistic incentives helping children to ask divergent questions can be very efficient. Divergent questions are questions that are relevant to the topic of a short text children read, but for which the answer is not in the text. At this point in time, a limit of these systems is that conversational agents are hand scripted. We would like in particular to equip them with capabilities to evaluate divergent questions children propose and give them feedback, and even better to generate examples of divergent questions given a new text. Thus, this project will study various approaches enabling this possibility.

Deep NLP context: The domain of natural language generation, and in particular question generation, has grown very fast in recent years, boosted by advances in learning very large language models such as GPT-3 (Brown et al., 2020). Several strands of works recently began to study the more specific challenge of generating automatically curiosity-driven questions (also called inquisitive questions, Ko et al., 2020), given a piece of text (Scialom et al., 2019), sometimes in the context of interaction with open-domain chatbots (Qi et al., 2020). One key challenge has been to build datasets to train question generation models, approached either through leveraging hyperlinked texts such as Wikipedia (Ferguson et al., 2020), or by crowdsourcing human annotations (Ko et al., 2020). However, these approaches are only first steps, and have neither considered adaptation to child cognition/educational domains, nor evaluation of the generation of curiosity-driven questions along dimensions that are fully relevant from linguistic and psychological perspectives.

Project:

The aim of this project will be to study how existing Deep NLP techniques for curiosity-driven question generation can be leveraged and adapted to a domain relevant for educational applications. In particular, the project will include:

1) Study and synthesis of the state-of-the-art in Deep NLP for curiosity-driven question generation, including experimenting/re-implementation of the most promising systems
2) Formalization of requirements for downstream educational applications, in strong interaction with senior researchers and PhD students in the team working on these educational applications

3) Building novel datasets for training and testing adapted to this domain (this may include collecting data using a crowdsourcing platform), associated with a scoring system grounded in the downstream educational application (e.g. evaluation along multiple dimensions including divergence, relevance, well-formedness, diversity of questions, etc).
4) Studying extensions of existing algorithms, or new algorithms, for generating diverse relevant curiosity-driven questions given a text, as well as for automatic scoring of questions proposed by children in the context of educational applications (or by the question generation system). Finetuning models pre-trained on existing adult-targeted language datasets will be considered.

Candidates will have the possibility to propose their own directions and ideas of approaches.

Requirements: We are looking for candidates with solid expertise in machine learning, especially deep learning and NLP, and associated software tools (pyTorch, tensorFlow, etc), as well as expertise in designing interactive web interfaces with javascript.

References

Abdelghani, R., Law, E., Oudeyer, P-Y., Sauzéon, H. (in review) Conversational agents for fostering curiosity-driven behaviors in children/

Alaimi, M., Law, E., Pantasdo, K. D., Oudeyer, P. Y., & Sauzeon, H. (2020, April). Pedagogical agents for fostering question-asking skills in children. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (pp. 1-13).

Becker, L., Basu, S., & Vanderwende, L. (2012, June). Mind the gap: learning to choose gaps for question generation. In Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (pp. 742-751).

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.

Ferguson, J., Gardner, M., Hajishirzi, H., Khot, T., & Dasigi, P. (2020). IIRC: A dataset of incomplete information reading comprehension questions. EMNLP 2020

Ko, W. J., Chen, T. Y., Huang, Y., Durrett, G., & Li, J. J. (2020). Inquisitive question generation for high level text comprehension. EMNLP 2020

Qi, P., Zhang, Y., & Manning, C. D. (2020). Stay hungry, stay focused: Generating informative and specific questions in information-seeking conversations. EMNLP 2020

Scialom, T., & Staiano, J. (2019). Ask to learn: A study on curiosity-driven question generation. arXiv preprint arXiv:1911.03350.